

RACF Data Management Policies

Introduction

The RHIC and Atlas Computing Facility (RACF) provides users with access to multiple types of data storage systems with different characteristics, i.e. service levels, performance, availability, etc.). These storage systems are available to users of the facility, be they individuals or collaborations, to help them accomplish their scientific research.

In order to provide current and future Project Principle Investigators with the necessary information to meet the Office of Science and Technology Policy and Department of Energy Office of Science data management requirements for long-term data sharing and preservation, this document provides the reference information on the RACF services in the area of data management. As a gross generalization, the RACF provides operational support for the data management capabilities that have been funded by individual projects. Long term access to individual project data, beyond the funded period is not guaranteed without a prior, written, service-level agreement.

Storage Resources

The storage resources provided by the RACF fall under the following categories

1. Transient data storage
2. Project dedicated data storage
3. Database services
4. Backup Services
5. HPSS Data Warehouse

Each category of storage resource contains several storage resources with differences in characteristics.

Transient data storage

The RACF operates two flavors of distributed file systems, AFS and NFS, for use by users that are available for use throughout the RACF facility. The NFS system is available for all users and provides storage space that is visible from all systems at the RACF. Contained within the NFS system are the user home directories. The primary use of NFS is to store files that are frequently read and written. With the exception of the home directories, all other areas within the NFS system are not backed up. Backups of the home directories are sent to the RACF backup system with a retention policy of 62 snapshots. However, all NFS data is stored on RAID 5/6 based storage systems with RAID controller and NFS server redundancy to maximize system availability and reduce the possibility of permanent

data loss due to hardware failures. Storage capacities within the NFS system are negotiated individually with each project and are commensurate with the funding allocated to or provided by the project to purchase and maintain the capacity. Outside of the home directory areas, allocation of storage, access controls and data retention within the NFS storage space is under the control of the individual projects. This includes policies on the deletion of files by people other than the owner of the files. The RACF may execute or enforce the data management policies for the individual projects, but the policies are determined by the projects (except in cases where there are Cybersecurity ramifications)

The AFS system is available for use by all users and provides storage space that is visible world- wide. The primary use of AFS is as storage and distribution resource for small files that are infrequently written or modified and read at moderate bandwidth from on-site and off-site locations. Unlike NFS, all AFS data is backed up to the RACF backup system. The retention period for backups of AFS data is 90 days. AFS data is stored on RAID 5/6 based storage systems with built-in RAID controller redundancy. AFS servers are all virtual machines on redundant physical hosts, with automatic virtual machine migration and start-up in event of a physical host failure.

Project dedicated data storage

The RACF facility operates data storage systems that are dedicated to individual projects, these data storage systems fall into two categories:

1. Compute node based storage
2. Dedicated storage node based storage

Regardless of the category, the availability, durability, and reliability of the data stored in these systems is individually negotiated with each project. In all cases, the storage systems are centrally located within the data center. Access, to these areas is also determined by negotiation with the individual projects. By default, there are no backups of these data stores. For the compute node based storage, individual compute nodes and the storage on the nodes are considered to be unreliable and subject to extended periods of unavailability. The compute node based storage consists of MapR based storage for the Astronomy group, Xrootd storage for STAR and US Atlas Tier 3, and dCache for PHENIX. In the case of the storage system with dedicated storage nodes, each individual node is considered to be reliable and is not subject to the availability constraints of the compute nodes.

Database Services

The RACF facility operates database servers that are dedicated to individual projects. Accessibility, availability, durability and reliability of the data and the database service are individually negotiated with each project. All database servers (and storage) are located within the data center.

Backup Services

The RACF operates two separate backup systems located on site, one dedicated to AFS backups and a second for other systems. By default, the retention period for backups is 90 days. There is no archival

storage within the two backup systems.

HPSS Data Warehouse

HPSS (High Performance Storage System) is a tape-based data storage system available to projects for long term, offline storage of bulk data. Direct access to HPSS is available via pftp, htar, and hsi to a restricted group of authorized users. Quality of service (bandwidth capabilities, storage capacities, data redundancy and access) are individually negotiated with projects and are commensurate with the funding provided by the projects. Since it is an automated tape library-based system, the actual performance is highly sensitive to I/O patterns.

External Data Access Capabilities

The RACF supports selected data access to storage systems from WAN (Wide Area Network) based clients. Direct WAN access to data located in NFS is provided through sftp gateways to any authorized user of the facility. Accessibility to individual files within the NFS system is determined by the identity of the sftp user AND the UNIX file and directory permissions on the directories. These permissions are set by the user that is assigned ownership of the directories and files. WAN transfers via sftp are encrypted by default and are therefore confidential over the wire. Direct WAN access to AFS-resident data is available to any user operating an OpenAFS client. Being a WAN accessible distributed file system, access controls to files within the AFS system are the responsibility of individual users, since they control the access control lists within the file system.

WAN access to the dedicated project data storage is negotiated individually with the projects and the service levels and mechanisms for access are determined by these negotiations. There is no WAN access to the HPSS data warehouse.

Internal Data Access Capabilities

Internally within the RACF, data is accessible from all systems having the respective AFS and NFS file systems mounted. Access to AFS is via the AFS protocol and access to NFS is via the NFS protocol. Access to the project-specific data stores is also available throughout the facility. Access protocols are defined by the available protocols supported by the storage system used by the specific data store and agreed upon by the individual projects and the RACF. Direct access to HPSS is available via pftp, htar, and hsi to a restricted group of authorized users.

Data Confidentiality and Access Controls

The RACF is a shared, open research facility intended for fundamental scientific research. Priority is given to facility access and sharing of information between users at the facility. Access controls in all cases is determined by the user. It is the user's responsibility to set file system permissions to set access controls according to their needs. With the exception of AFS, all identities are determined through UNIX user and group ID's on individual systems, this includes HPSS and the project-dedicated data storage systems.

RACF system administrators with "root" or "administrator" privileges are not constrained by file system permissions and have the ability to bypass all access control mechanisms. "root" privileges are

used only under certain highly restricted situations, and typically only look at user files when there is a problem or a security issue. In the former case, this is typically in response to a request from owner of the file.